

UNCLASSIFIED

Defense Technical Information Center
Compilation Part Notice

ADP011121

TITLE: Thrust Vector Control and Visualisation for Stovl Aircraft

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Active Control Technology for Enhanced Performance Operational Capabilities of Military Aircraft, Land Vehicles and Sea Vehicles
[Technologies des systemes a commandes actives pour l'amelioration des performances operationnelles des aeronefs militaires, des vehicules terrestres et des vehicules maritimes]

To order the complete compilation report, use: ADA395700

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:

ADP011101 thru ADP011178

UNCLASSIFIED

Thrust Vector Control and Visualisation for Stovl Aircraft

P.M. Lodge & C. Fielding

BAE SYSTEMS

Warton Aerodrome

Lancashire, PR4 1AX, UK

Summary

There are many areas where the design of control laws for short take-off and vertical landing (STOVL) aircraft are radically different from those of a conventional design. One the most challenging areas is related to control of the thrust vector, and in particular, how to design control laws that can deal with conflicting demands, control saturation and integrator wind-up. Unlike a conventional aircraft, where full primary control surface deflections are rarely used, the STOVL aircraft nozzle vector angle and engine thrust are often operated on their limits for significant periods of time. For example, to achieve maximum deceleration from wing-borne flight, the nozzles are vectored to their forward authority limits and remain there for several seconds.

The first part of this paper provides background on the UK's Integrated Flight and Propulsion System (IFPCS) programme, the BAE SYSTEMS P112C-6 aircraft configuration and the Rolls-Royce RB571 engine concept. A command strategy to control the aircraft during wing-borne flight and through the transition to the hover is introduced, which leads to the description of a suitable control law architecture to satisfy this command strategy.

With the approach chosen, 'thrust vector equations' are required to transform the pilot's commands into nozzle vector angle and engine thrust demands. The equations are derived for the Harrier and subsequently for the more complex P112C-6 configuration. These equations are then extended to cover the cases of saturated control conditions. Emphasis is placed on the visualisation and verification of the complex functions that result.

Discussion of the application of the thrust vector equations in a non-linear real-time simulation, including flight and engine control system integration aspects, is also covered and the paper concludes with lessons learned in this area.

Introduction

For over 30 years, BAE SYSTEMS has been involved in the design and development of control concepts for future STOVL aircraft. The majority of this research has been undertaken in support of the UK's VAAC Harrier [1,2] and IFPCS [3,4] programmes and has been aimed at improving the understanding of aircraft handling and flight control for transition and jet-borne flight.

One of the most interesting and challenging areas of this research has been in the design of control laws and in particular, development of thrust vector control strategies. This has resulted from the desire to reduce pilot workload in the launch

and recovery of the aircraft [5], where direct open-loop control of throttle and nozzle angle can be replaced by closed-loop control. This allows the pilot to command the variables of primary interest and importantly, in the reference axes that are most appropriate to the flight regime and task.

Experience has shown that the algorithms required to implement the desired thrust vector control strategy can become complex and difficult to interpret and verify, without adequate visualisation tools. This paper will demonstrate a thrust vector control strategy appropriate to future STOVL aircraft, with the emphasis on the visualisation, verification and understanding of the functions required for implementation.

The UK's IFPCS Programme

The UK's Integrated Flight and Propulsion Control System (IFPCS) Technical Demonstration Programme is a joint UK MoD / Industry-funded research effort, which has mainly focused on investigating aircraft control and system architectures to meet the challenges associated with future STOVL aircraft.

The project has utilised the P112C-6 canard-delta aircraft configuration (shown below), which is largely based on BAE SYSTEMS' experience with the EAP aircraft. The main configuration differences are that the P112C-6 has side intakes and is a close-coupled canard-delta, as opposed to the chin intake and long-coupled canard-delta of the EAP aircraft.

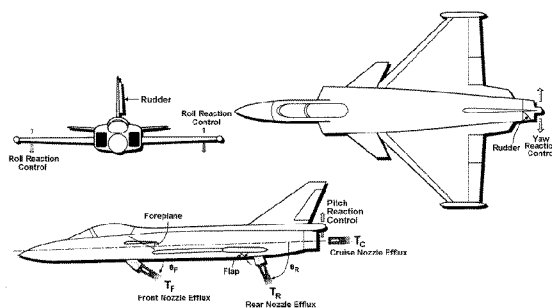


Figure 1. P112C-6 aircraft general arrangement

The project powerplant is the Rolls-Royce RB571-10 direct-lift engine concept, which has two distinct modes of operation:

- Flight mode – in which the lift system is disabled and the engine operates as a conventional turbo-fan as shown below.

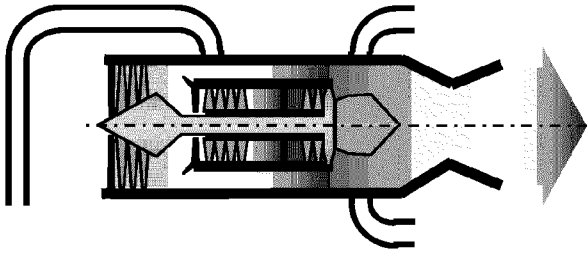


Figure 2a. RB571-10 engine configured in Flight mode

- Lift mode – in which the main core flow is directed through two rear-mounted vectoring nozzles and a proportion of the by-pass air is ducted through to a front lift nozzle. This provides a three-poster remotely unaugmented lift system.

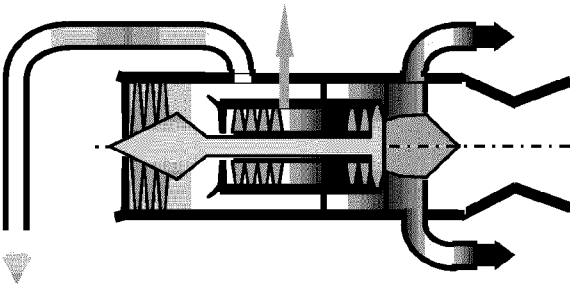


Figure 2b. RB571-10 engine configured in Lift mode

In Lift mode, the engine configuration allows the ratio between the front and rear thrust components to be varied, providing a source of pitching moment control in the hover. The front and rear nozzles are independently actuated, effectively providing four independent motivators:

- engine gross thrust, X_G
- engine differential thrust, X_S
- front nozzle angle, θ_F
- rear nozzle angle, θ_R

A three-axis reaction control system, using engine bleed-air is also available.

In order to effect a transition between the Lift and Flight modes, the engine must undertake a mode change. This entails opening and closing the internal doors that enable (or disable) the lift system, and reconfiguring the engine control system.

Control strategy and architecture

In order to significantly reduce pilot workload whilst providing acceptable handling qualities in recovering the aircraft to a ship or dispersed site, the decision was made to adopt the 'two-inceptor' longitudinal control strategy shown below. This allows the pilot to control flight path using the right hand, and speed using the left hand, throughout the transition from wing-borne flight to the hover. Between 100 and 70 knots, there is a blending region, within which, the pilot's commands become referenced to earth-axes.

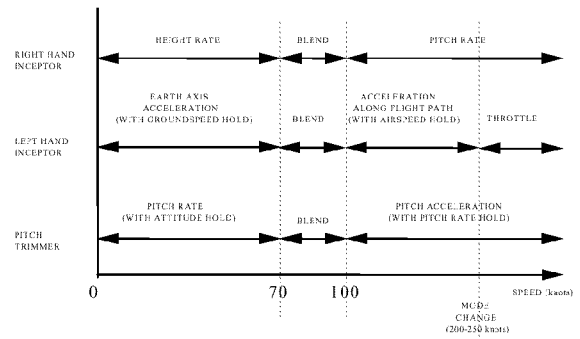


Figure 3a. Longitudinal control strategy

The lateral / directional control strategy adopted is shown below and is fairly conventional. However, the blend between bank angle and roll rate commands is worthy of note. This blend could be placed over a lower speed range, but has been chosen to allow the pilot to command a rate in each axis (height rate and turn rate), thereby providing control harmonisation.

Automatic turn co-ordination in pitch is also provided, throughout the Lift mode. This has been shown to be a significant contributor to workload reduction.

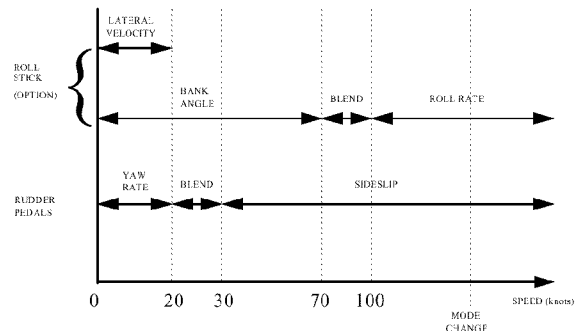


Figure 3b. Lateral / directional Control Strategy

In adopting this approach, high levels of closed-loop augmentation are required in each axis. This can only be realistically provided through a full authority digital flight control system. A typical high-level control law architecture is shown below.

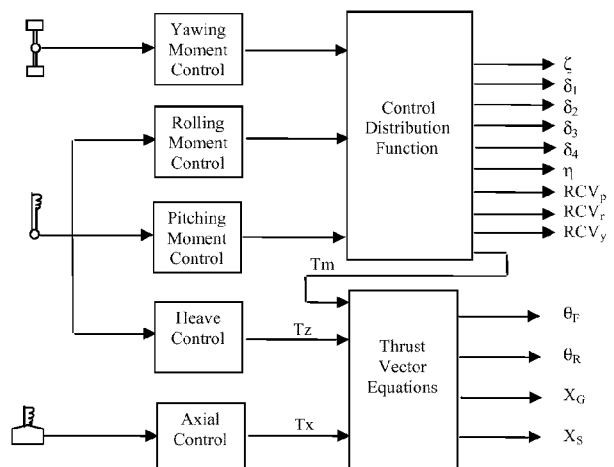


Figure 4. Typical control law architecture

The control distribution function routes the moment controller outputs to the appropriate effector(s), taking into account control redundancy and the variation in control effectiveness across the flight envelope. It can route the pitching moment demand to either the flaps, foreplane, pitch RCS or the thrust vector equations.

Where the dynamic pressure is high enough, the aircraft is controlled conventionally via aerodynamic forces and moments. However, as dynamic pressure is reduced, the vectored thrust forces and moments (T_x , T_z and T_m) are phased in, eventually completely replacing the aerodynamic terms in the hover condition. The Thrust Vector Equations are used to transform the thrust vector force and moment demands (T_x , T_z and T_m), into front and rear nozzle angle and thrust component demands.

The Harrier Configuration

Prior to analysing the P112C-6 configuration in detail, it is worth considering the Harrier aircraft configuration. This is a relatively simple arrangement (shown below) based around a four-post lift system.

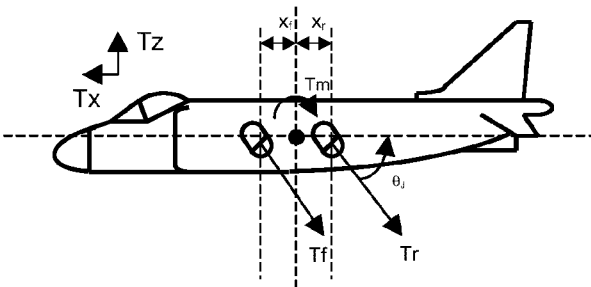


Figure 5. Harrier lift system

The front and rear lift-posts are provided by exhausting air from the engine fan and turbine stages, respectively, through two pairs of ducted nozzles, located fore and aft of the centre of gravity. Movement of the aircraft's centre of gravity is restricted such that thrust-induced pitching moments are minimised, but also that any pitching moment generated by the engine thrust is well within the trimming capability of the aircraft's reaction control system. This feature helps to make the aircraft flyable during the transition and hover, with a minimum of control augmentation.

Introducing automatic control to such a configuration by including the thrust vector within the feedback loops can introduce several new problems. For example, thrust vector saturation, in terms of dealing with conflicting demands and control law integrator conditioning. This aspect has been successfully addressed for the Harrier aircraft through the DERA's VAAC Harrier research programme [1], in which two-inceptor control laws were flight tested. However, the solutions to these problems are fairly straightforward, due to the relative simplicity of the Harrier's lift system.

The increased levels of augmentation needed to provide a two-inceptor control strategy result in functional complexity, in order to automate the management of the thrust vector. There are several ways that this might be done, such as the use of Trimmaps [6] or by using Non-linear Dynamic Inversion [7]. Here, the basic vectored thrust equations for the Harrier are represented by a simple static inversion of the equations used to transform the body-axis rectangular components of demanded thrust (T_x and T_z) into nozzle and throttle demands:

$$\text{Nozzle angle } (\theta_j) = \tan^{-1}(T_z / T_x)$$

$$\text{Throttle } (T_f + T_r) = K \sqrt{(T_x^2 + T_z^2)}$$

It should be noted that at low speed, where there is little or no aerodynamic control power, control of the aircraft's pitching moment is performed by the reaction control system alone.

The above equations are only valid if there is adequate nozzle angle and throttle quadrant ranges to satisfy the demands. When this is not the case, and either the throttle or nozzle angle demand reach an authority limit (i.e. become saturated), modified control action must be taken.

The diagram below shows a 2-D plot of a strategy for dealing with saturation of the thrust vector for the Harrier configuration.

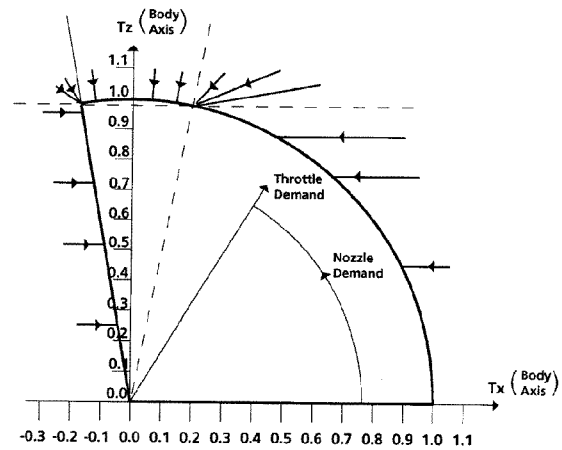


Figure 6. Saturation strategy for the Harrier

The diagram shows how demands that exceed the capability of the thrust vector are mapped back onto the available throttle and nozzle boundaries. It can be seen that generally, excessive T_x demands are mapped back onto the maximum throttle arc, by satisfying the T_z demand and reducing the T_x demand. This approach assigns priority to the vertical axis and ensures that heave transients are minimised.

Difficulty can be encountered in ensuring that the algorithms required to provide the above strategy are continuous and cover all possible combinations of pilot/FCS demands. In order to provide confidence that the algorithms perform their intended role, a pictorial representation of the operation of the algorithms is useful. Such a representation is shown in the figures below where advantage is taken of 3-D graphics. It should be noted that the T_x and T_z demands have been non-dimensionalised.

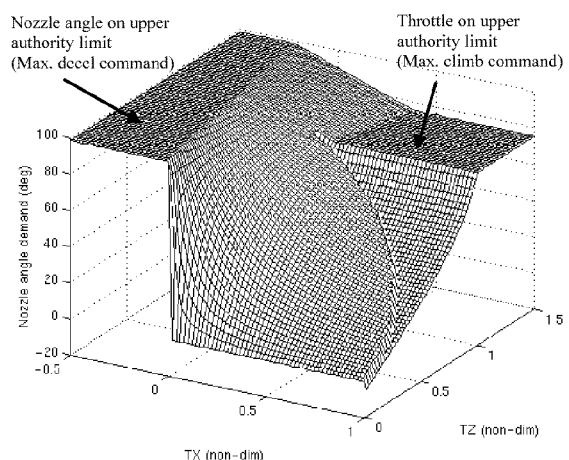


Figure 7a. Harrier nozzle angle visualisation

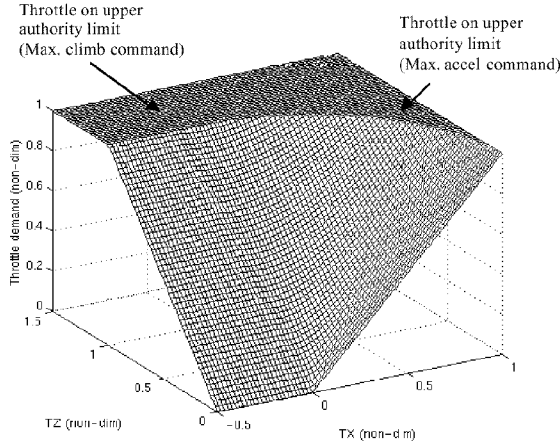


Figure 7b. Harrier throttle visualisation

Thrust vector equations for P112C-6

Having established the thrust vector equations for the simpler Harrier configuration, the more complex P112C-6 configuration can now be examined.

The figure below shows the P112C-6 aircraft, with its lift system geometry marked up with x- and z-axis distances between the centre of gravity and the points at which the front and rear thrust vector components act.

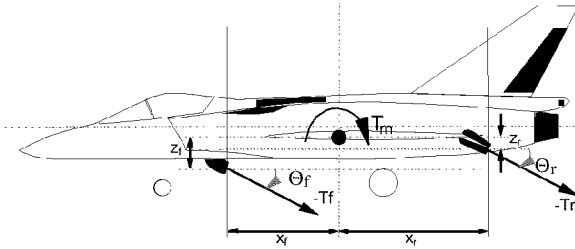


Figure 8. P112C-6 Lift System Geometry

By considering the above diagram, a static relationship can be derived between the front and rear thrust vector components and the body axis longitudinal forces and moment.

$$T_x = T_f \cos(\theta_f) + T_r \cos(\theta_r)$$

$$T_z = -T_f \sin(\theta_f) - T_r \sin(\theta_r)$$

$$T_m = T_f (X_f \sin(\theta_f) + Z_f \cos(\theta_f)) + T_r (Z_r \cos(\theta_r) - X_r \sin(\theta_r))$$

Since there are four control variables but only three longitudinal degrees of freedom, there exists a level of redundancy. By expressing the front lift nozzle vector angle as a function of T_z and T_x , as in the Harrier example above:

$$\theta_f = \tan^{-1} (T_z / T_x)$$

the above equations can be re-arranged to give three thrust vector equations that satisfy the force and moment demands.

$$\tan(\theta_r) = \frac{T_m + T_z \cot(\theta_f) (Z_f - Z_r) + T_z X_f - T_x Z_r}{\cot(\theta_f) (T_m - T_x Z_f - T_z X_r) - T_x (X_f + X_r)}$$

$$T_f = \frac{T_m - T_x Z_r - T_z X_r}{(X_f + X_r) \sin(\theta_f) + (Z_f - Z_r) \cos(\theta_f)}$$

$$T_r = \frac{(T_m - T_x Z_f \cos(\theta_f) - T_x X_f \sin(\theta_f))}{(Z_r - Z_f) \cos(\theta_r) \cos(\theta_f) - X_r \sin(\theta_r) \cos(\theta_f) - X_f \sin(\theta_f) \cos(\theta_r)}$$

It should be noted that the above equations are unbounded, and are valid only where there exists adequate control authority to satisfy the force and moment demands. Upon saturation of one or more of the control effectors, in order to prevent any stability and control problems that would almost certainly ensue, the thrust vector equations must be modified to maintain control of the aircraft.

It is important to assign priority to each of the three degrees-of-freedom to determine in which order the control loops are to be 'opened' in the event of a first saturation.

Since the pilot is not directly commanding nozzle angle and throttle position, it is likely that when maximum commands are input, the resulting force demands will exceed the available performance capabilities of the aircraft. In this situation, at least one of the demands must be reduced to allow priority to be assigned to another.

Experience has shown that the highest priority should be allocated to the pitching moment demand in order to ensure that attitude control is maintained. The 'long-coupled' P112C-6 thrust vector geometry exhibits the ability to generate very large pitching moments when configured in its Lift mode, which can far exceed the pitch RCS authority. These moments must be controlled to within acceptable limits at all times. Allocating the highest priority to pitching moment control provides two benefits.

- It minimises uncommanded pitch interaction, improving the perceived handling.
- In extreme circumstances it helps provide a stable platform for ejection.

Medium priority should be allocated to the vertical thrust component to ensure that no significant uncommanded height change can occur.

Lowest priority should be allocated to the horizontal thrust component. This is most likely to result in acceleration/deceleration performance limitations, but can in certain circumstances (for example in low speed tight turns) result in an uncommanded forward speed change.

By monitoring the front and rear nozzle angles and thrust components against their respective limits, thrust vector saturation can be detected and corrective action taken within the control laws.

These principles are best demonstrated by means of an example. For instance, if the pilot commands a straight-in maximum deceleration, it results in the front and rear nozzle angles moving to their forward authority limits. By setting the rear nozzle angle demand to the forward limit (to generate maximum negative T_x) and neglecting the T_x demand, equations can be obtained for T_f and T_r , in terms of T_m , T_z , θ_f and θ_r :

$$T_r = \frac{\sin(\theta_f) (-T_z X_f - T_m) - T_z Z_f \cos(\theta_f)}{\sin(\theta_f) (X_f \sin(\theta_r) - Z_r \cos(\theta_r) + X_r \sin(\theta_r)) + Z_f \cos(\theta_f) \sin(\theta_r)}$$

$$T_f = \frac{-(T_r \sin(\theta_r) + T_z)}{\sin(\theta_f)}$$

By performing this process for each single saturation and then again for each double saturation, modified thrust vector equations result, which are switched in and out, depending upon the level of saturation and the combination of control devices involved.

As with the Harrier example, difficulty can be encountered in ensuring that the algorithms and switching logic that constitute

the vectored thrust equations, are continuous and cover all possible combinations of demands, without excessive transients. With the increased complexity of the P112C-6 geometry, it is important to gain confidence by verifying that the implementation is performing as intended. For example, for a maximum deceleration command, one would expect the rear nozzle angle to be demanded to its forward limit. For a maximum climb command, one would expect the front and rear nozzle thrusts to be demanded to their maximum values.

The complexities of the trigonometric functions and the switching logic provide plenty of scope for an incorrect functional implementation. It is extremely difficult to identify such errors from the coding, since detailed knowledge of how the compiler deals with trigonometric functions might be required. Tabulation of the solutions of the equations is helpful, but produces a mass of information that is difficult to check. It has been found to be beneficial to plot this tabulated data as a series of 3-D plots. The figures below depict the saturation effects, where a nose-up pitching moment of 15 kNm is required from the propulsion system. A series of plots can be produced by varying the pitching moment required.

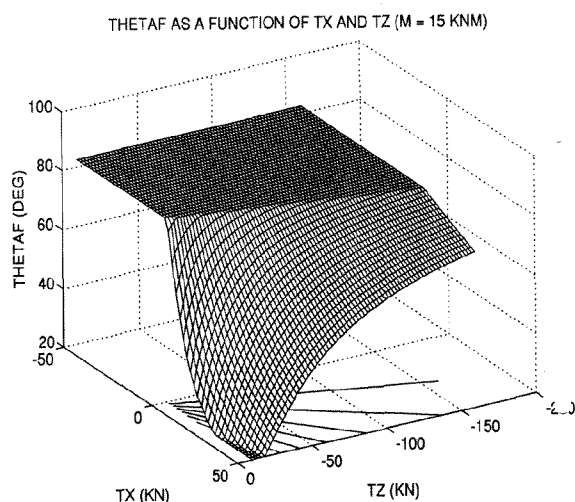


Figure 9a. P112C-6 front nozzle angle visualisation

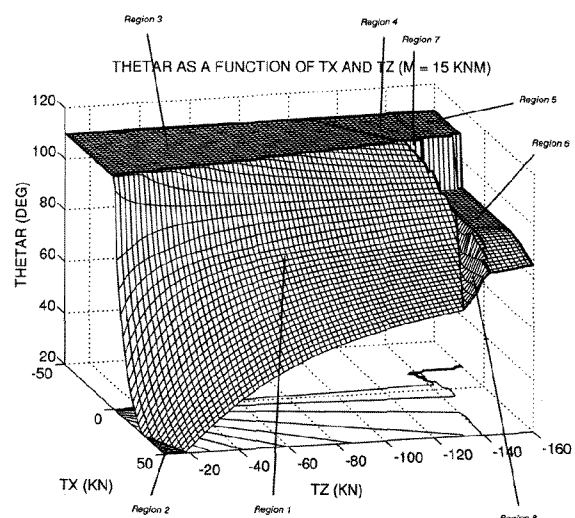


Figure 9b. P112C-6 rear nozzle angle visualisation

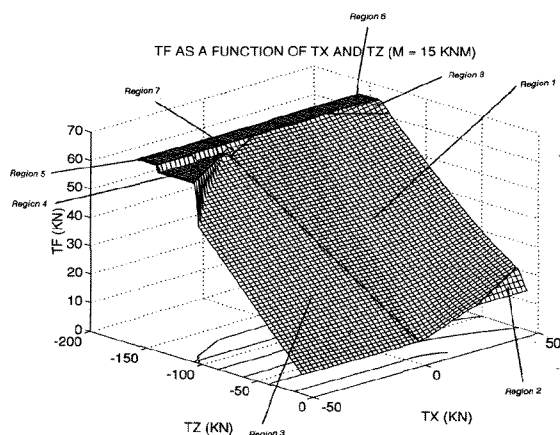


Figure 9c. P112C-6 front thrust visualisation

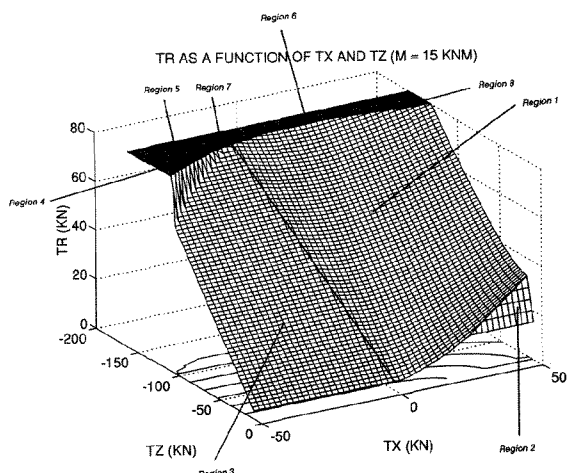


Figure 9d. P112C-6 rear thrust visualisation

Each of the marked up regions on the plots corresponds to a saturated condition:

- Region 1 – Unsaturated – not control limited.
- Region 2 – Single saturation - θ_R lower authority limit (maximum deceleration / maximum sink commanded).
- Region 3 – Single saturation - θ_R upper authority limit (maximum deceleration commanded).
- Region 4 – Double saturation - T_R & θ_R upper authority limits (maximum deceleration / maximum climb commanded).
- Region 5 – Triple saturation - T_R , T_F , θ_R upper authority limits (maximum deceleration / maximum climb commanded).
- Region 6 – Double saturation - T_R & T_F upper authority limits (maximum climb commanded).
- Region 7 – Single saturation - T_R upper authority limit (maximum climb commanded).
- Region 8 – Single saturation - T_F upper authority limit (maximum climb commanded).

This form of plot has proved to be invaluable in the development of the thrust vector equation algorithms and for verifying the implementation. It also aids the understanding of the control of the aircraft.

Development of the thrust vector equations

Ultimately, the thrust vector equations will form a major portion of the functional integration of the flight and propulsion control systems. In this role, the thrust vector equations will reside in the flight control computer where they will be required to take signals from the propulsion control system. These signals will indicate saturation of the engine controls, both in terms of gross thrust and thrust split.

In developing the thrust vector equations it is advisable to employ a staged development process:

- The first stage has been covered in the previous section, where the algorithms and switching logic that constitute the thrust vector equations have been fully tested, in an open-loop manner, over the full range of operating conditions.
- The equations need to be linearised and appropriate feedback loops designed to control the aircraft motions. The action of the controllers on the error signals will provide the T_x , T_z and T_m demands that need to be satisfied.
- In the next stage, the thrust vector equations need to be implemented in a closed-loop simulation of the aircraft, including a model of the engine and the full control laws. At this stage, the engine model should consist of a simple low order transfer function representation with notional, engine gross thrust and thrust split limits, as this will ensure rapid development and allow confidence to be gained in the closed-loop operation.
- Having established a satisfactory closed-loop arrangement, increased engine complexity can be introduced. At this stage, signals need to be synthesised from the engine control system variables, which can be fed into the thrust vector equations, to indicate when engine saturation has occurred.

Integrator anti-wind-up

Integral action is generally used in each of the control loops, to minimise drift in the commanded variable. In this case, specific action is required to prevent the integrators from 'winding-up' in the presence of saturation.

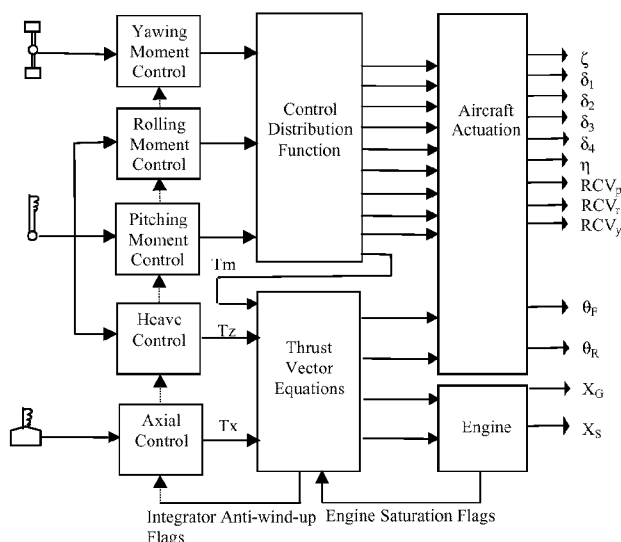


Figure 10. Functional arrangement of integrator anti-wind-up logic

A suitable control law structure is shown above, including the three separate control loops for the Tx, Tz and Tm demands, the thrust vector equations and the signals used to freeze the integrators. Also shown are the feedback signals from the propulsion control system (PCS) that are used to monitor engine saturation.

Simulation development and visualisation

At this stage, full development testing can start. Initially, off-line analysis using a full-envelope simulation model of the aircraft, engine and control system should be used to develop the control laws to an adequate standard for pilot-in-the-loop simulation. This model should include:

- A full aerodynamic representation of the aircraft including jet-induced effects. This should cover the extreme values of incidence and sideslip that can be achieved at low speed.
- A full non-linear representation of the flight control system, including control laws and hardware.
- A full non-linear representation of the engine and its control system.

During pilot-in-the-loop simulation, the operation of the engine and nozzles should be observed and recorded, as this will give an early warning, should a thrust vector-related problem occur. A typical engineering console display is shown below. Included on this display are all the aircraft and engine control devices.

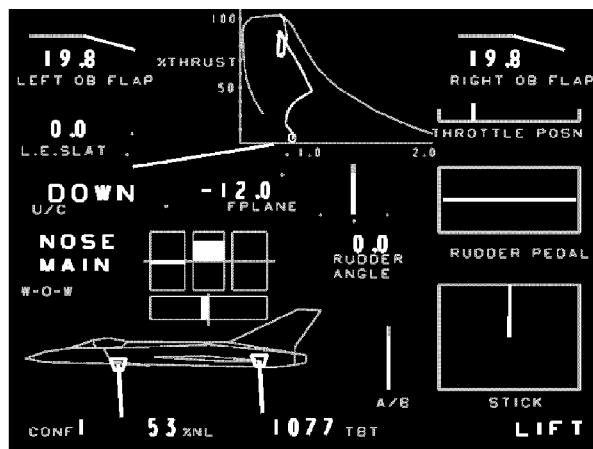


Figure 11. Simulator console engineering display

The engine response is represented by a 2-D plot of gross thrust versus thrust ratio (top middle). Experience has shown that a locus plot of the engine response is useful for observing the engine behaviour in relation to its control boundaries. Typically, the engine response locus should remain visible for between 2 and 5 seconds, to prevent excessive clutter. The front and rear nozzle vector angles are shown on the aircraft representation, which also moves with aircraft pitch angle, giving an earth-referenced view. Each of the aerodynamic surfaces is represented through a simple hinged surface and the RCS valve openings are represented by simple rectangles, representing the linear opening of each of the RCS valves.

Closed-loop response characteristics

The figure below shows the aircraft response to a full stick pull (+5 m/s height rate command) in the hover.

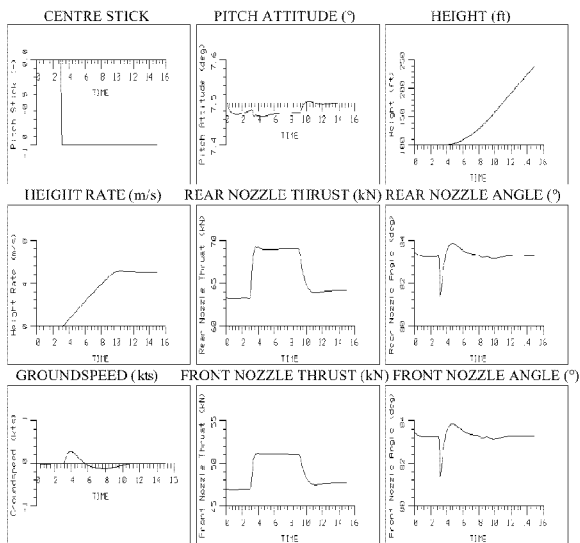


Figure 12a. Response to full back stick in the hover

It can be seen that the front and rear thrust limits are saturated for about 6 seconds, whilst the aircraft height rate builds up to the commanded value. The excellent level of decoupling can be seen in the minimal disturbances in pitch attitude ($\leq \pm 0.1^\circ$) and groundspeed ($\leq \pm 0.5$ kts). The vectored thrust equations and anti-wind-up logic are working as intended.

The figure below shows the aircraft response to a maximum deceleration command (-0.4 g axial acceleration) at 80 knots.

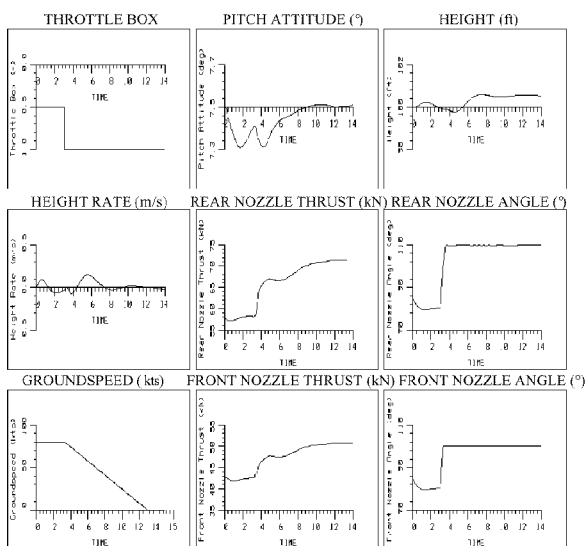


Figure 12b. Response to maximum deceleration at 80 knots.

It can be seen that the front and rear nozzle angles reach their respective forward limits in response to the command and remain there throughout the response. Again the excellent level of decoupling can be seen in the minimal disturbances in height ($\leq \pm 1$ ft) and pitch attitude ($\leq \pm 0.2^\circ$). Once again, the thrust vector equations are working as expected.

Whilst these responses show excellent decoupling characteristics, they are for the baseline model, where the thrust vector geometry is accurately measured. Studies have been performed to gain insight into how the decoupling degrades with centre-of-gravity movement. This showed that quite large movements could be tolerated (± 0.2 m), whilst maintaining satisfactory decoupling. The effect of varying aircraft mass is

also a significant effect, as this changes the thrust margin. However, the integral action in each of the control loops helps to minimise decoupling and drift.

Lessons learned

In developing the thrust vector equations for the P112C-6 configuration, a number of lessons have been learned, which are listed below:

- It is considered essential to perform a staged development of the thrust vector equations, in order to gain confidence in their operation at each stage.
- The algorithms required to implement the thrust vector equations contain complex trigonometric functions and switching logic. Visualisation of these functions is vital in building confidence in their operation.
- An advantage of the analytical expressions used in the thrust vector equations, is that they can be easily differentiated to provide analytical expressions for the local gradients. These can then be used in the linear design of the feedback loops. This also eliminates the requirement to store a large amount of linearised data.
- Engine model complexity should be considered at each stage of the development. Generally, two or three engine representations will exist, each exhibiting various levels of complexity. These may range from a simple second-order transfer function representation, with rate and amplitude limiting, to a full thermodynamic representation of the engine with its control system. Experience has shown that much progress can be made early on in development with a very simple engine representation.
- The dynamic response characteristics of the engine / airframe combination are key to providing acceptable handling. Undesirable response characteristics in either can result in design difficulties. It is therefore important that the engine and airframe development teams work together effectively to enable a balanced design.
- Saturation of the thrust vector can be minimised by careful limiting of the pilot's commands and the Tx and Tz demands, upstream of the thrust vector equations. This effectively limits the accelerations required of the aircraft and should be achieved after considering the aircraft performance specification and the impact on aircraft handling. For example, minimising the Tx and Tz demands to improve decoupling at the expense of aircraft performance should be considered, but only if the performance penalty is marginal and acceptable.

Discussion

The concept of the vectored thrust equations, as presented in this paper, are a static inversion of the forces and moments associated with an aircraft's lift system. The equations form part of the flight control laws, in conjunction with other inverse terms such as those required to normalise aerodynamic control powers (to correct for changes in dynamic pressure). Dynamic conversion terms are excluded here, but could be added, to ensure common phasing of all the thrust vector components. This would help to minimise the coupling due to out-of-phase forces and moments.

The scheme described in this paper is just one of a number of options available for designing control laws for a complex STOVL aircraft. Other related approaches to control law design,

such as the use of Trimmings [6] or Non-linear dynamic inversion [7], can be used to develop a satisfactory set of control laws.

All three approaches can potentially produce extremely good results for a datum aircraft configuration. There is a further major challenge in ensuring that the control system provides adequate handling over the full range of aircraft configurations and over the full flight envelope. For methods employing full non-linear dynamic inversion, there is the potential for an extremely complex inversion task, which may lead to adoption of on-board aircraft and engine models, with their associated complexity.

The complexity of the aircraft configuration and desired operating envelope will help to decide on the level of inversion that is necessary to generate the control system design. It is expected that all approaches will encounter similar challenges, since these will be mainly associated with the physical layout of the aircraft configuration and its propulsion system. For example, some of the major challenges include the determination of the following:

- The most appropriate pilot command strategy.
- The most effective use of the aerodynamic and propulsion controls.
- Satisfactory decoupling of aircraft responses.
- Correct conditioning of integrators and moding logic.
- Satisfactory handling qualities for transition and jet-borne flight. This is the subject of ongoing research [8].

This paper has outlined some of the aspects to be considered when designing flight control laws for a STOVL aircraft. The theme of the paper has been the visualisation of some of the control law functions, as this is believed to be an important aspect of any advanced design, which is complex and very non-linear. The use of advanced graphics should be used wherever possible to aid the understanding and to support the verification of complex systems. It is believed that this is an area of future research that may be beneficial for future systems, which will no doubt, further increase in complexity, in line with advances in computing capability.

Acknowledgements

The authors would like to acknowledge the support of the UK Ministry of Defence Procurement Agency and the Defence Research Agency for funding STOVL research programmes that provided the background to this paper.

References

1. C.Fielding & S.L.Gale, 'Flight Demonstration of an Advanced Pitch Control law in the VAAC Harrier Aircraft', AGARD-CP-560, 1995.
2. G.T.Shanks, S.L.Gale, C.Fielding & D.V.Griffith, 'Flight Control and Handling Research with the VAAC Harrier Aircraft', in TISCHLER,M.B. (Ed.): 'Advances in Aircraft Flight Control', (Taylor and Francis, 1996).
3. D.A.Allen & R.J.Sleeman, 'The Integrated Flight and Propulsion Control System Programme', International Powered Lift Conference, London, 1994.
4. C.Fielding, ' Design of Integrated Flight and Propulsion Control Systems', AGARD Conference Proceedings, CP-548, 1993.

5. J.D.B.Paines, 'Optimising STOVL Flight Control Functionality for Flying Qualities', International Powered Lift Conference, London, 1998.
6. B.W.Rawnsley, S.J.Andrews, G.W.D'Mello, L.D.Gillett 'The Application of Trimmings to the Design of STOVL Aircraft', International Powered Lift Conference, London, 1998.
7. J.Franklin, 'Revised Simulation Model of the Control System, Displays, and Propulsion System for an ASTOVL Lift Fan Aircraft', NASA Tech. Memo 112208, 1997.
8. J.C.Gibson, P.M.Lodge & C.Fielding, 'Development of handling qualities criteria for STOVL Aircraft', Powered Lift Conference, Arlington, 2000.

Paper #26

Q by Daniel Walker: Are the complex phenomena (e.g. engine behavior) modeled well enough that they (i.e. the models) provide useful information over-and-above the small perturbation model used in the control system design?

Are there any nasty dynamics that make this a potentially tricky contest problem?

A. (P. M. Lodge): Yes, the small perturbation model is derived directly from a full non-linear thermodynamic model of the engine and its control system. The interaction of the engine thrust flow field with the airframe is also fully modeled.

Yes, there is much potential for coupled response from the engine producing heave/pitching moment transients. The flight control laws can be used to minimize the coupling, but it is preferable for the FCS and PCS designers to work together, understand each other's problems and develop a balanced solution.

This page has been deliberately left blank



Page intentionnellement blanche